

Deep Learning based Segmentation of Fish in Noisy Forward Looking MBES Images

Jesper Haahr Christensen^{***} Lars Valdemar Mogensen^{**}
Ole Ravn^{*}

^{*} *Electrical Engineering Department, Technical University of Denmark,
2800 Kgs. Lyngby, Denmark,
e-mail: {jehchr,or}@elektro.dtu.dk*

^{**} *ATLAS MARIDAN, 2960 Rungsted Kyst, Denmark,
e-mail: {jhc,lvm}@atlasmaridan.com*

Abstract: In this work, we investigate a Deep Learning (DL) approach to fish segmentation in a small dataset of noisy low-resolution images generated by a forward-looking multibeam echosounder (MBES). We build on recent advances in DL and Convolutional Neural Networks (CNNs) for semantic segmentation and demonstrate an end-to-end approach for a fish/non-fish probability prediction for all range-azimuth positions projected by an imaging sonar. We use self-collected datasets from the Danish Sound and the Faroe Islands to train and test our model and present techniques to obtain satisfying performance and generalization even with a low-volume dataset. We show that our model proves the desired performance and has learned to harness the importance of semantic context and take this into account to separate noise and non-targets from real targets. Furthermore, we present techniques to deploy models on low-cost embedded platforms to obtain higher performance fit for edge environments – where compute and power are restricted by size/cost – for testing and prototyping.

Keywords: Autonomous Underwater Vehicle (AUV), Deep Learning, Semantic Segmentation, Sonar Imaging, Multibeam Echosounder (MBES) Imaging, Fish Monitoring.

1. INTRODUCTION

In many circumstances, the preferred choice of an imaging device is an optical sensor. This produces high quality, high-resolution images with plenty of information on detailed content, colors, shapes, and textures. However, in certain environments, optical sensors are not well suited. In underwater imaging, optics are highly affected by turbidity in water which reduces visibility. This is both in terms of obscuring line of sight with particles and other organic detritus (marine snow) and in terms of illumination, which is attenuated with the amount of turbidity and depth.

As such, a multibeam echosounder (MBES) is often preferred for underwater imaging. In terms of imaging information, it hardly compares with optics; however, the robustness to underwater conditions and extended range while maintaining a somewhat high resolution of contours makes MBES ideal. Even compact sonars integrable on small autonomous underwater vehicles (AUVs) or remotely operated vehicles (ROVs) offer ranges up to more than 100m with range and angular resolution dependent on application, model, and price. Sonar images do, however, suffer from distortion, noise and contain low-level detail and visual information only and do as such still pose a challenging task when processing sonar data for vision purposes.

In certain marine applications such as subsea monitoring, inspection tasks, or fisheries, it may be useful to have a

small easy-portable system capable of simultaneously collecting and utilizing sonar data for autonomy or meta-data purposes. By using semantic-based segmentation masks, fish instances can be identified and measured in size and numbers. Fish schools can be identified at range, and absolute position can automatically be mapped, or relative position can feed into a navigation system. For databases with meta-data generated by such a model, queries can be made to specific situations and thus save hours and days of labor-intensive manual work of reviewing and labeling sonar data in extensive datasets.

The architectural choices for our segmentation model leverage recent advances in Deep Learning (DL) and Convolutional Neural Networks (CNNs) within segmentation for medical imaging (Ronneberger et al. (2015)) and scene understanding (Long et al. (2014); Badrinarayanan et al. (2017); Chen et al. (2018)). Our proposed model allows for end-to-end training and processing from input sonar image to binary output mask.

As this work is highly application-oriented, we further demonstrate how we prepare our model for deployment in an edge environment, i.e., on-board online processing right where the data is created. Here size, power, and cost often limit the options for compute. We present tools to accomplish this and report performance on two typically used low-cost embedded platforms.

The contributions in this work comprise a novel application proposal to obtain a target/non-target probability

prediction for all range-azimuth positions as projected by a MBES on a pixel-based coordinate space using DL on the edge. Along with model architecture, techniques for training and obtaining high accuracy models with low data volumes are demonstrated. We also note that the proposed work does not limit itself within the application of fish segmentation but could potentially be expanded to any target or multi-class segmentation models.

The remainder of this paper is organized as follows. Section 2 provides a short literature survey of related work. Section 3 describes the datasets and the collection of these. Section 4 describes our method for fish segmentation. Section 5 provides details of experiments. Section 6 demonstrates deployment and performance on the edge. Finally, in Section 7, we present our conclusion.

2. RELATED WORK

Deep Learning In line with CNNs revolutionizing the area of visual perception above water, increasing efforts are being made to apply such methods for underwater optical and acoustical sensing as well. One key challenge for applying such methods in the maritime domain is data availability. For autonomous driving and related fields, large annotated datasets containing millions of images are openly available, such as ImageNet (Russakovsky et al. (2014)), Kitti (Geiger et al. (2012)), BDD100K (Yu et al. (2018)) and Cityscapes (Cordts et al. (2016)) to name a few. Sadly this is not yet the case in our field and thus in the case of a small dataset (<5.000 samples) techniques such as data augmentation (Perez and Wang (2017)) and transfer learning (Yosinski et al. (2014)) may be used. In Christensen et al. (2018), an object detection model, using both before mentioned techniques, has been trained to detect, localize and classify fish and fish species in optical images. Valdenegro-Toro (2017); Fuchs et al. (2018) applies transfer learning to perform image classification on specific cropped regions of sonar images. Similarly, Horimoto et al. (2018) train an object detection model to predict the presence and location of sea turtles in sonar images.

Sonar Image Segmentation In Sture et al. (2018), a naïve form of segmentation is achieved on high-resolution synthetic aperture sonar images for mapping corals in the data. This is carried out by dividing sonar trajectories into smaller segments and then perform a per-segment classification and reassemble the map to get a pseudo segmentation of the traversed area. Dos Santos et al.



Fig. 1. Photo of the ATLAS MARIDAN man-portable AUV used for collecting training data.

(2017) obtain a segmentation of sonar images by first identifying blobs or regions of interest in the image using traditional image processing techniques. Second, detected blobs are classified using linear machine learning models such as support vector machines, decision trees, and K-nearest neighbors.

3. DATA

We introduce two separate datasets for training and testing our model. The data is collected using the same sensor but with two different vehicles at different locations and fish species. Although the visual appearance of the two datasets is similar, completely separated datasets for training and testing are a prerequisite in order to validate the generalization of the model rather than only testing the fit to the training data.

3.1 Training data details

Our training set is collected autonomously in the Danish Sound off the coast of our offices in Rungsted (DK). This is carried out using our man-portable AUV shown in Fig. 1. The MBES data is collected using a forward-looking BluePrint Oculus m750d multibeam imaging sonar. Data is recorded with a range interval of $[0, 20]$ m, an aperture of 130° horizontal and 20° vertical at an operating frequency of 750 kHz.

The labeled training set contains in total 50 images, which include more than 5000 targets (Herring) along with non-fish data such as surface reflections, surface




Fig. 2. Training set MBES data samples: (a) school of Herring, (b) surface vessel, (c) school of Herring.

vessels, and bottom returns. The dataset is annotated using the LabelBox¹ annotation tool. Three samples from the training set are shown in Fig. 2. Far from all fish in every image have been annotated due to the extensive process of manually drawing segmentation masks.

3.2 Test data details

Our test set is collected in the fjords of the Faroe Islands. This is carried out using our ROV sensor test rig shown in Fig. 3(a). The sensor and settings are similar to those described in Section 3.1. The recordings differ with varying range settings, which scales the visual appearance. The targets of the Faroe Islands are of the fish specie Wittling, which slightly differs from Herring in size, numbers, and behavior. The setup and a sample are shown in Fig. 3.

4. METHOD

Our network architecture is shown in Fig. 4 and builds upon typical convolutional encoder-decoder style networks. Here we embed our input in a latent feature space and then seek to reconstruct the input as semantic classes. We employ skip connections to fully recover the fine-grained spatial information lost in pooling or down-sampling layers. This is as proposed by Long et al. (2014) and used in nearly all following work on segmentation. Table 1 and the following describes the architecture and dimensions in more detail.


Input dimensions are $1 \times 320 \times 128$ corresponding to the 1-channel MBES data at half resolution represented in a Cartesian coordinate space. As most targets are not more than a few pixels in height, this is the borderline resolution without losing too much information.

Each **Conv Layer** combines a sequence of a convolution, batch normalization, dropout, convolution, and batch normalization. This seeks to obtain an efficient and regularized training progress. The ReLU non-linearity is used as activation function for each **Conv Layer**. Each convolution is furthermore padded in order to keep input and output dimensions consistent.

¹ <https://labelbox.com/>



(a)



(b)



(c)

Fig. 3. Test data collection at the Faroe Islands: (a) Our ROV test rig for sensors, (b) Sonar image sample, (c) Optical image sample.




Fig. 4. The encoder-decoder style network architecture for fish segmentation in sonar images. Input/output resolution is 320×128 and lowest resolution (in the bottleneck) is 20×8 .

Table 1. Architecture of fish segmentation network.

	Name	Feat maps (input)	Feat maps (output)
Encoding path	Conv Layer 1	$1 \times 320 \times 128$	$16 \times 320 \times 128$
	Max-pooling 1	$16 \times 320 \times 128$	$16 \times 160 \times 64$
	Conv Layer 2	$16 \times 160 \times 64$	$32 \times 160 \times 64$
	Max-pooling 2	$32 \times 160 \times 64$	$32 \times 80 \times 32$
	Conv Layer 3	$32 \times 80 \times 32$	$64 \times 80 \times 32$
	Max-pooling 3	$64 \times 80 \times 32$	$64 \times 40 \times 16$
	Conv Layer 4	$64 \times 40 \times 16$	$128 \times 40 \times 16$
	Max-pooling 4	$128 \times 40 \times 16$	$128 \times 20 \times 8$
	Bottleneck	$128 \times 20 \times 8$	$256 \times 20 \times 8$
Decoding path	Up-sample 1	$256 \times 20 \times 8$	$256 \times 40 \times 16$
	Conv Layer 5	$256 \times 40 \times 16$	$128 \times 40 \times 16$
	Up-sample 2	$128 \times 40 \times 16$	$128 \times 80 \times 32$
	Conv Layer 6	$128 \times 80 \times 32$	$64 \times 80 \times 32$
	Up-sample 3	$64 \times 80 \times 32$	$64 \times 160 \times 64$
	Conv Layer 7	$64 \times 160 \times 64$	$32 \times 160 \times 64$
	Up-sample 4	$32 \times 160 \times 64$	$32 \times 320 \times 128$
	Conv Layer 8	$32 \times 320 \times 128$	$16 \times 320 \times 128$
	Sigmoid Layer	$16 \times 320 \times 128$	$1 \times 320 \times 128$

Up-sample combine transpose convolutions and a fusion with the corresponding feature map in the encoding path.

Finally, a **Sigmoid Layer** is applied. This combines a 1×1 convolution to down-sample the channel dimension to a single output channel and a sigmoid activation to reach the final output as a probability for each pixel. The sigmoid output is then threshold at 0.5 to produce either 1 or 0 as the final output for the two segmentation classes; fish and non-fish.

Since we have two classes only (fish, non-fish), we can learn the mapping from input sonar images to output masks using binary cross-entropy loss:

$$\begin{aligned} \mathcal{L}_{\text{BCE}}(M) &= - \sum_{i=1}^{C=2} y_i \log [M(x_i)] \\ &= -y_1 \log [M(x_1)] - (1 - y_1) \log [1 - M(x_1)] \end{aligned} \quad (1)$$

where M is our model, C is our two classes, i represents each class in C , y is the ground truth pixel value, and x is the input data.

Our objective therefore is:

$$\min_M \mathcal{L}_{\text{BCE}}(M). \quad (2)$$

5. EXPERIMENTS

The model is trained using one NVIDIA Tesla K80 GPU. As the architecture is fairly light-weight in terms of parameters, time to reach convergence is less than 1 h. The model is implemented in the TensorFlow DL framework using the Keras frontend for Python. During training, we use a batch size of 4 and all weights are initialized using the He normal initializer (He et al. (2015)). All ReLUs are leaky with slope 0.2. The loss function is minimized using the recently introduced RAdam solver (Liu et al. (2019)) with initial learning rate $\eta = 0.5 \times 10^{-4}$ and parameters β_1 and β_2 set to 0.9 and 0.999 respectively.

With these hyperparameters, the model converges within 100 epochs to a validation loss and an accuracy of 0.059 and 97%, respectively. Additionally, the precision, recall,

Table 2. Edge performance.

	Default (non-optimized)	TFLite (RPi3)	TFLite (RPi4)	TensorRT (Nano)
FPS	0.37 / 0.76 / 11	0.45	1.43	33
Speedup	-	1.2x	1.88x	3x
Model size	22.5 MB	2.4 MB	2.4 MB	11.5 MB
Size reduction	-	9.375x	9.375x	1.96x

and F1-score for the fish-class (the content we are interested in) are 69%, 81%, and 75%, respectively. The final model size is 22 MB.

To overcome issues with low data volume, the data is heavily augmented at training time. This is done by randomly applying horizontal and vertical flipping, rotations in range $[-20, 20]^\circ$, width and height shifts in range $[-20, 20]\%$, and random cropping.

We test our model on the dataset obtained in the fjords of the Faroe Islands. This allows for a qualitative assessment only, as we have no ground-truth data due to the labor-intensive work going into creating this. It will, however, indicate the performance, robustness, and generalization achieved by the model. Fig. 6 shows nine samples of sonar images and the predicted output mask from the trained model. The samples clearly show that the model has learned to discriminate between fish and non-fish and thus predicts noise, surface reflections, bottom returns, and the like as non-fish.

As explained in Section 3.2, the targets in the test data are no longer Herring, but a bigger fish specie named Wittling. The visual (sonar image) appearance of these fish is similar at an instance-level. However, the schooling behavior is unlike Herring, and for much of the obtained data Wittlings are notably more scattered and at a closer range than Herring in the training data.

6. EDGE IMPLEMENTATION

Operating at the edge refers to a setting where data is processed at the location of the sensor. Typically for small submersible systems, the budget for cost, compute, power, and size for processing platforms is restricted. Two popular choices are a Raspberry Pi and an NVIDIA Jetson Nano. Table 2 summarizes performance measures on post-training optimized models on a single CPU core on a Raspberry Pi and on an NVIDIA Jetson Nano using its embedded GPU. Details are presented in the following.

As this work is highly application-oriented and part of ongoing work at ATLAS MARIDAN, the success of a model relies not only on results achieved of the model but also on its ability to be running in an embedded environment on the edge. Here we demonstrate how models quick and straightforward can be prepared for such environments while still maintaining a Python environment for convenient prototyping during inference. We report performance on two common and popular embedded boards without requirements for external accelerators such as the Intel Movidius Neural Compute Stick or the Google Coral USB Accelerator.

Raspberry Pi 3/4 Using the TensorFlow framework, we first save the model as a static inference graph and




Fig. 6. Test images from the Faroe Islands dataset. Sonar image is on top, predicted mask on bottom.

second quantize the weights of the graph to 8-bit precision rather than 64-bit. The quantization is performed using the TensorFlow Lite framework for on-device inference. The post-quantization model size is 2.4 MB and does inference at 1.3 FPS using a single core only (out of 4).

This is a more than $\times 9$ reduction in model size and a nearly $\times 2$ speedup.

NVIDIA Jetson Nano Using the high-performance inference platform TensorRT from NVIDIA, we can use in-build support in TensorFlow to convert our model to a

post-training quantized 16-bit TensorRT graph. As above, we first save our model as a static inference graph and then convert the graph to a quantized TensorRT graph.

The TensorRT model size is 11.5 MB and does inference at 33 FPS using the embedded GPU on the Jetson Nano. This is a $\times 2$ reduction in model size and a $\times 3$ speedup.

7. CONCLUSION

Our model converges on the validation dataset with 97% accuracy. The precision, recall, and F1-score for fish targets are 69%, 81%, and 75%, respectively. Additionally, the results presented from the test dataset in Fig. 6 clearly show that our model has learned the desired functionality of segmenting fish-like targets from noise, surface reflections, and other non-fish objects in sonar images. Some targets remain to be recognized, and the model may be affected by the low data volume or bias caused by incomplete annotations. Possible improvements could be achieved by retraining with new data processed by the model in a semi-supervised manner.

From the performance summary in Table 2, it is shown that the model is capable of performing inference at suitable processing times on low-cost embedded devices. We obtain 33 FPS on a Jetson Nano and more than 1 FPS using only one core (out of 4), and nearly no memory on a Raspberry Pi 4. 1 FPS is deemed sufficient for most monitoring and long-range tracking applications at ATLAS MARIDAN.

REFERENCES

- Badrinarayanan, V., Kendall, A., and Cipolla, R. (2017). Segnet: A deep convolutional encoder-decoder architecture for image segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(12), 2481–2495. doi:10.1109/TPAMI.2016.2644615.
- Chen, L., Papandreou, G., Kokkinos, I., Murphy, K., and Yuille, A.L. (2018). Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40(4), 834–848. doi:10.1109/TPAMI.2017.2699184.
- Christensen, J.H., Mogensen, L.V., Galeazzi, R., and Andersen, J.C. (2018). Detection, localization and classification of fish and fish species in poor conditions using convolutional neural networks. In *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, 1–6. doi:10.1109/AUV.2018.8729798.
- Cordts, M., Omran, M., Ramos, S., Rehfeld, T., Enzweiler, M., Benenson, R., Franke, U., Roth, S., and Schiele, B. (2016). The cityscapes dataset for semantic urban scene understanding. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Dos Santos, M., Ribeiro, P., Núñez, P., Drews-Jr, P., and Botelho, S. (2017). Object classification in semi structured environment using forward-looking sonar. *Sensors*, 17(10), 2235. doi:10.3390/s17102235. URL <http://dx.doi.org/10.3390/s17102235>.
- Fuchs, L.R., Gällström, A., and Folkesson, J. (2018). Object recognition in forward looking sonar images using transfer learning. In *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, 1–6. doi:10.1109/AUV.2018.8729686.
- Geiger, A., Lenz, P., and Urtasun, R. (2012). Are we ready for autonomous driving? the kitti vision benchmark suite. In *Conference on Computer Vision and Pattern Recognition (CVPR)*.
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. *IEEE International Conference on Computer Vision (ICCV 2015)*, 1502. doi:10.1109/ICCV.2015.123.
- Horimoto, H., Maki, T., Kofuji, K., and Ishihara, T. (2018). Autonomous sea turtle detection using multi-beam imaging sonar: Toward autonomous tracking. In *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, 1–4. doi:10.1109/AUV.2018.8729723.
- Liu, L., Jiang, H., He, P., Chen, W., Liu, X., Gao, J., and Han, J. (2019). On the variance of the adaptive learning rate and beyond.
- Long, J., Shelhamer, E., and Darrell, T. (2014). Fully convolutional networks for semantic segmentation. *CoRR*, abs/1411.4038. URL <http://arxiv.org/abs/1411.4038>.
- Perez, L. and Wang, J. (2017). The effectiveness of data augmentation in image classification using deep learning. *CoRR*, abs/1712.04621. URL <http://arxiv.org/abs/1712.04621>.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In N. Navab, J. Hornegger, W.M. Wells, and A.F. Frangi (eds.), *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, 234–241. Springer International Publishing, Cham.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M.S., Berg, A.C., and Li, F. (2014). Imagenet large scale visual recognition challenge. *CoRR*, abs/1409.0575. URL <http://arxiv.org/abs/1409.0575>.
- Sture, Ø., Ludvigsen, M., Scheide, M.S., and Thorsnes, T. (2018). Recognition of cold-water corals in synthetic aperture sonar imagery. In *2018 IEEE/OES Autonomous Underwater Vehicle Workshop (AUV)*, 1–6. doi:10.1109/AUV.2018.8729718.
- Valdenegro-Toro, M. (2017). Best practices in convolutional networks for forward-looking sonar image recognition. *CoRR*, abs/1709.02601. URL <http://arxiv.org/abs/1709.02601>.
- Yosinski, J., Clune, J., Bengio, Y., and Lipson, H. (2014). How transferable are features in deep neural networks? In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2, NIPS'14*, 3320–3328. MIT Press, Cambridge, MA, USA. URL <http://dl.acm.org/citation.cfm?id=2969033.2969197>.
- Yu, F., Xian, W., Chen, Y., Liu, F., Liao, M., Madhavan, V., and Darrell, T. (2018). BDD100K: A diverse driving video database with scalable annotation tooling. *CoRR*, abs/1805.04687. URL <http://arxiv.org/abs/1805.04687>.